

IN THE SPECIFICATION:

Please find attached a series of amendments to certain of the pages to show corrected material.

At page 12 beginning at line 24, (and continuing through page 13), change the paragraph to read as follows:

With reference to FIG. 2, a flowchart illustrating the different phases performed to diagnose the cluster is shown. The process begins with start bubble 31 followed by a process step (block 32) to carry out the initialization phase. The process continues with an inquiry as to whether or not the initialization phase (See Fig. 3) set a starting flag to the installation form (diamond 33). If the answer to this inquiry is yes, the installation phase is carried out (block 34), as indicated in Fig. 3. If the answer to this step 33 inquiry is no, the installation phase is not carried out (see Fig. 4). The process continues with a process step (block 35) to carry out the diagnostics phase, as indicated in Figs. 5A, 5B, 5C, 5D, 5E. [[Fig. 4.]] After the diagnostics phase, the process continues by carrying out the results phase (block 36) shown in Fig. 2. [[5).]] The process then ends (bubble 37).

At page 13, line 13, change the paragraph to read as follows:

Referring now to FIG. 3, a flowchart illustrating the steps for the initialization phase is shown. The process begins with a start bubble 41 followed by an inquiry as to whether or not the node contains previously stored data (diamond 42). The node in this step is the current node. If the answer to this inquiry is no, the process sets a starting flag to the installation form (block 47) and exits at (bubble 50). Otherwise, on YES at step 42, the process continues by gathering data (at step 43) previously stored for the current node (block 43). After gathering this data, the process continues with an inquiry as to whether or not the node contains current data (diamond 44). If the answer to this inquiry is no, the process sets a starting flag to the installation form (block 47) and exits (bubble 50). If the answer to this inquiry at step 44 is yes, the process gathers data about the current state of the node (block 45). The process continues with an inquiry as to whether or not there is a discrepancy between the current data and the previously stored data (diamond 46). For example, if the IP Address for the cluster was saved as 192.59.200.18 and the currently derived IP Address was now 192.59.200.180, this would indicate a discrepancy between the two configurations. If the answer to this inquiry is yes, the process sets the starting flag to the installation form (block 47) and exits (bubble 50). Otherwise, the process at step 46 (NO) continues with an inquiry as to whether or not the node contains companion node data (diamond 48). A companion node is another node known within the cluster distinct from the current node. If the answer to this inquiry is yes, the process gathers data about the companion nodes (block 49) and exits (bubble 50). If the answer to this inquiry is no, the process ends (bubble 50).

Beginning at page 14 (and continuing through page 15, lines 4-6), change the paragraph to read as follows:

With reference to FIG. 4, a flowchart that illustrates the steps for the installation phase is shown. The process begins with start bubble 61 followed by a process step (block 62) to display the installation form. The installation form allows the user to input information that cannot be derived directly from the node. Data for the node, both derived and manually input, appears on categorized tabs for clarity. The process continues with a process step (block 63) to allow the user to change appropriate data for the node. For instance the user could set the IP Address for the cluster. Next, the process compares the data, including the information for operation of the node, network communications, and shared disk storage, for the node as defined on the form with the established guidelines, such as those found in the Microsoft Cluster Support Knowledge Base for use with MSCS, for clustering this node (block 64). After comparing, the process reports any discrepancies to the user (block 65). The process continues with an inquiry as to whether or not another node for the cluster is known (diamond 66). If the answer to this inquiry is no, the process ends (bubble 72). If the answer to this inquiry is yes, the process specifies the other node as the companion node for determining the ability of the two nodes to form a cluster (block 67). Next, the process displays any discrepancies in data, such as a difference in the IP address for the cluster, to the user (block 68). This is followed by an inquiry as to whether or not serious discrepancies[[,]] (such as the IP addresses used to communicate between the two nodes being on different subnets)[[,]] exist between the two nodes (diamond 69). If the answer to this inquiry is yes, the process allows the user to correct

discrepancies (block 70). If the answer to this inquiry is no, the process does not allow for the correction of discrepancies. The process continues with a process step (block 71) to save configuration data. The configuration data will contain a set of categories dealing with the configuration of the cluster. Each category will contain specific data items related to the category along with the value defined for the aspect. For instance, the Cluster category will have a data item for the IP Address for the cluster and will have a value such as 192.59.200.18. The process then ends (bubble 72).

At page 15, lines 27-28, change the paragraph to read as follows:

Referring now to FIG. 5A, the diagnostics phase begins with start bubble 81 and continues with a process step (block 82) of displaying the diagnostics form. The diagnostics form will specify and control the level, type, and execution of the diagnostics to perform. Next, the process allows the user to specify (from a list of nodes known to the cluster) a companion node, a node distinct from the current node that the cluster uses, for determining the ability of the nodes to act as a cluster (block 83). This companion node will hold for the diagnostics tests during the current execution of the program until re-specified. The process continues with a process step (block 84) to allow the user to set the test level for the diagnostics used during the session. The test level can be set to one of two levels, A non-obtrusive test disallows any diagnostics that will compromise the execution of the node. The complete test level allows any diagnostic to be used. After the test level is set using option buttons, the process allows the user to specify the test categories also found on the diagnostics form to use during the session (block 85). These categories contain a collection of related tests that the user may select. The process continues by allowing the user to set option buttons to indicate the method for interaction of the program with the user (block 86). Depending on the selection of the method for interaction, the process may either run all tests regardless of outcome, or stop when a test reports an error. The interaction process may also require the user to step between tests or allow

the process to simply continue without user intervention. The interaction process also allows the user to specify the number of loops to make through the tests. The process then continues as described in FIG. 5B.

At page 18, line 7, change the paragraph to read as follows:

With reference to FIG. 5C, the Diagnostic Phase process continues with the diagnostics phase step 94 by inquiring whether or not the testing of communications is selected. These tests will diagnose physical problems communicating between the nodes within the cluster. If the answer to this inquiry is no, the process continues as described in FIG. 5E via markers C and E of Fig. 5D. Otherwise, at step 94 (YES) an inquiry is made as to whether or not testing of connections marked as private is selected (diamond step 95). The private communications pathway will be used for communications between nodes in a cluster to determine their accessibility. If the answer to this inquiry is yes, the process step 96 tests the ability for the current node to communicate with the companion node via all network connections marked as private (block 96). Such a connection performs only communications between the clustered nodes used to control the cluster. Any errors are reported to the user (block 97). If the answer to the inquiry posed by diamond step 95 is no, the process does not perform this test. The diagnostic process continues with an inquiry as to whether or not testing connections marked as public is selected (diamond 98). The public communications pathway will be used for communications between a node and other servers external to the cluster. If the answer to this step 98 inquiry is yes, the process tests the ability for the current node to communicate with the companion node via all network connections marked as public (block 99). Such a connection performs normal communications not associated with the

operation of the cluster. The process then reports any errors to the user (block 100). If the answer to the inquiry posed by diamond step 98 is no, the process does not perform this test. The process then continues as described in FIG. 5D, via marker D.

At page 20, line 10, change the paragraph to read as follows:

Next, an inquiry is made at step 107 as to whether or not the testing of command execution is selected. Command execution allows the current node to initiate commands on the companion node. If the answer to this inquiry is yes, the process tests the ability of the current node to execute commands on the companion node (block 108). Such a connection will allow the operation of the clustering software (17 and 20 [[25]] of Fig. 1) to be effective. Any errors are reported to the user (block 109). If the answer to the inquiry in diamond 107 is no, the process does not perform this test. The process then continues as specified in FIG. 5E via the marker E.

At page 22, line 24, change the paragraph to read as follows:

Described herein has been a method for re-constituting a multiple mode server system cluster after a cluster failure has occurred. The re-establishment of the clustered server system is enabled by a series of method steps utilizing an initialization phase, an installation phase, a diagnostics phase, and a results phase which indicates to the user how the failure can be corrected to re-establish the server cluster into proper operation.